Interprocedural Type Specialization of JavaScript Programs Without Type Analysis

Maxime Chevalier-Boisvert
DIRO, Université de Montréal, Quebec, Canada
chevalma@iro.umontreal.ca

Marc Feeley
DIRO, Université de Montréal, Quebec, Canada
feeley@iro.umontreal.ca

arXiv:1511.02956v1 [cs.PL] 10 Nov 2015

Categories and Subject Descriptors D.3.4 [*Programming Languages*]: Processors—compilers, optimization, code generation, run-time environments

Keywords Just-In-Time Compilation, Dynamic Language, Optimization, Object Oriented, JavaScript

Abstract

Dynamically typed programming languages such as Python and JavaScript defer type checking to run time. VM implementations can improve performance by eliminating redundant dynamic type checks. However, type inference analyses are often costly and involve tradeoffs between compilation time and resulting precision. This has lead to the creation of increasingly complex multi-tiered VM architectures.

Lazy basic block versioning is a simple JIT compilation technique which effectively removes redundant type checks from critical code paths. This novel approach lazily generates type-specialized versions of basic blocks on-the-fly while propagating context-dependent type information. This approach does not require the use of costly program analyses, is not restricted by the precision limitations of traditional type analyses.

This paper extends lazy basic block versioning to propagate type information interprocedurally, across function call boundaries. Our implementation in a JavaScript JIT compiler shows that across 26 benchmarks, interprocedural basic block versioning eliminates more type tag tests on average than what is achievable with static type analysis without resorting to code transformations. On average, 94.3% of type tag tests are eliminated, yielding speedups of up to 56%. We also show that our implementation is able to outperform Truffle/JS on several benchmarks, both in terms of execution time and compilation time.

[Copyright notice will appear here once 'preprint' option is removed.]

1. Introduction

The highly dynamic semantics of JavaScript (JS) make optimization difficult. Late binding, dynamic code loading and the eval construct make type analysis a hard problem. Precise type analyses also tend to be expensive, and are often considered too costly to be used in Just-In-Time (JIT) compilers.

Lazy Basic Block Versioning (BBV) [12] is an intraprocedural JIT compilation strategy which allows rapid and effective generation of type-specialized machine code on-the-fly without a separate type analysis pass (Section 2.2).

In this paper we introduce an interprocedural variant which extends BBV with mechanisms to propagate type information interprocedurally, across function call boundaries (Section 3), both function parameter and return value types. Combining these new elements with BBV yields a lightweight approach to interprocedurally type-specialize programs on-the-fly without performing global type inference or type analysis in a separate pass.

A detailed evaluation of the performance implications is provided in Section 5. Empirical results across 26 benchmarks show that, on average, the interprocedural BBV approach, as far as eliminating type tests goes, performs better than a simulated perfect static type analysis (Section 5.1). On average, 94.3% of dynamic type tests are eliminated. Speedups of up to 56% are also obtained, with only a modest increase in compilation time.

2. Background

The work presented in this paper is implemented in a research Virtual Machine (VM) for JavaScript (ECMAScript 5) known as Higgs[1]. The Higgs VM features a JIT compiler built upon an experimental design centered around BBV. This compiler is intended to be lightweight with a simple implementation. Code generation and type specialization are performed in a single pass. Register allocation is done us-

[1] https://github.com/higgsjs

```
function f(n) {
    if (n == 0)
        return 0;
    else
        return n + f(n-1);
}
```

Figure 1. Simple example of lazy basic block versioning.

ing a greedy register allocator. The runtime and standard libraries are self-hosted, written in an extended dialect of ECMAScript with low-level primitives. These low-level primitives are special instructions which allow expressing type tests as well as integer and floating point machine instructions in the source language.

2.1 Value Types

Higgs segregates values into categories based on type tags [18]. These type tags form a simple, first-degree notion of types that is used to drive code versioning. The `unknown` type is also used by the code versioning process to indicate that the type is unknown. These type identifiers are listed in the following table.

`int32`	32-bit integers
`float64`	64-bit floating point values
`null`	JS `null`
`const`	miscellaneous JS constants
`string`	JS strings
`array`	JS arrays
`closure`	JS function objects
`object`	other JS objects
`unknown`	type is unknown

2.2 Lazy Basic Block Versioning

BBV is a JIT code generation technique originally applied to JavaScript by Chevalier-Boisvert & Feeley [12], and adapted to Scheme by Saleil & Feeley [27]. The technique bears similarities to HHVM's tracelet-based compilation approach and Psyco's JIT code specialization system [26].

Consider the recursive JavaScript function `f` shown in Figure 1. Given that `f`'s parameter `n` is not known to be an `int32`, `float64`, or other type, a non-optimizing JIT compiler would put several type tests in the generated code, namely for the operators `==` (type of `n`), `-` (type of `n`), and `+` (type of `n` and type of `f(n-1)`). With intraprocedural BBV a JIT compiler would determine the type of `n` at the first type test (in `n==0`), and, assuming it is an `int32`, would generate code for the expression `n+f(n-1)` specialized to `n` being an `int32`, thus avoiding the two subsequent type tests of `n`.

Note that intraprocedural BBV does not propagate type information between caller and callee functions and so two type tests remain: for the expression `n==0` because the type of `f`'s parameter is unknown, and for the addition, because `f`'s return value type is unknown. This is the motivation for our work.

The BBV approach efficiently generates type-specialized machine code in a single pass, without the use of costly type inference analyses or profiling. This is achieved by lazily cloning and type-specializing single-entry single-exit basic blocks on-the-fly. As in Psyco, code generation and code execution are interleaved so that run-time type information can be extracted by the code generation engine. The accumulated information allows the removal of redundant type tests, particularly in performance-critical paths.

Type information is accumulated as blocks and type tests are compiled, and propagated forward to successor blocks. A mapping of live variable types to specialized block versions is maintained. The technique is a form of forward type propagation. It is unlike traditional fixed point type analyses in that instead of computing a fixed point on value types, it is lazily computing a fixed point on the creation of basic block versions. The presentation in [12] explains strategies to effectively limit the generation of block versions and prevent a code size explosion. In practice a hard limit of 5 versions per basic block yields good results.

Because of its JIT nature, BBV has at least two powerful advantages over traditional static type analyses. The first is that BBV focuses on the parts of the control flow graph that get executed, and it knows precisely which they are, as versions are only generated for executed basic blocks. The second is that code paths can often be duplicated and specialized based on different type combinations, making it possible to avoid the loss of precision caused by control flow merges in traditional type analyses.

2.3 Typed Object Shapes

BBV, as presented in [12], deals with function parameter and local variable types, and it has no mechanism for attaching types to object properties. This is problematic because, in JS, functions are typically stored in objects. This includes object methods and also global functions (JS stores global functions as properties of the *global object*). To allow interprocedural type propagation it is essential to know which function is being called for as many call sites as possible, both for calls to global functions and method calls.

Currently, all commercial JS engines have a notion of object shapes, which is similar to the notion of property maps invented for the Self VM. That is, any given object contains a pointer to a shape descriptor providing its memory layout: the properties it contains, the memory offset each property is stored at, as well as attribute flags (i.e. writable, enumerable, etc.). Work done on the Truffle Object Model [29] describes how object shapes can be straightforwardly extended to also encode type tags for object properties, so long as property writes are guarded to update object shapes when a property type changes.

Chevalier-Boisvert and Feeley have extended upon the original BBV work with typed object shapes [11], giving

```
function Accum() {
    this.n = 0;
    this.add = function id1(x) { this.n += x };
    this.sub = function id2(x) { this.n -= x };
}

var a = new Accum();
a.add(5);
```

Figure 2. Accumulator object with two methods.

Figure 3. Shape of object a, encoding the identity of both of its methods.

BBV the ability to propagate information about the shape of objects, and also their property types. Polymorphic Inline Caches (PICs) are broken down into cascades of shape tests exposed as multiple basic blocks in the compiler Intermediate Representation (IR). This makes it possible for the BBV engine to extract and propagate shape information at every property read and write. Propagating said information then allows extracting property types from object shapes at property reads, but also to eliminate redundant PICs after successive property reads and writes.

An interesting advantage of typed shapes is that the identity of methods, when known, is also encoded in object shapes. In practice, this makes it possible to determine the identity of callees in the large majority of cases. This is done by extending the type `closure` so that it carries the identity of the method.

Consider, for instance, the JS program given in Figure 2. The identity of the methods assigned in the `Accum` function are encoded in the shape of the newly constructed object and propagated to the variable `a`. Hence, in the expression `a.add`, the variable `a` is known to have the shape given in Figure 3, and the method call to `a.add` is known to refer to the `id1` method at code generation time.

In Higgs, method identity is encoded as a unique pointer to the IR node of that method. It is denoted `closure/id`N in the examples. There is also the type `closure/unknown` needed to handle the situation where it is known that a value is a function, but the identity of the function is unknown (this gives more precision than using the `unknown` type, which is used when nothing is known of the type). When writing a value of type T to property P of an object with shape S, if T is different than $S.P$ (the type of P in the shape S), the compiler will generate code to create and store in the object a new shape S' identical to S except that $S'.P = T$.

Method identity is valuable information because it makes it possible to directly jump to a callee's entry point without dynamic dispatch. The usefulness of this in the implementation of interprocedural BBV is explained further in Section 3.1.

3. Interprocedural Lazy BBV

Here we present the main contributions of this paper: entry point specialization and call continuation specialization.

3.1 Entry Point Specialization

Procedure cloning has been shown to be a viable optimization technique, both in ahead of time and JIT compilation contexts. By specializing function bodies based on argument types at call sites, it becomes possible to infer the types of a large proportion of local variables, allowing effective elimination of type checks.

Our first extension to BBV is to allow functions to have multiple type-specialized entry points. That is, when the identity of a callee at a given call site is known at compilation time, the JIT compiler requests a specialized entry point for the callee which assumes the argument types known at the call site. Type information is thus propagated from the caller to the callee.

Inside the callee, BBV proceeds as described in [12], deducing local variable types and eliminating redundant type checks. Our approach places a hard limit on the number of versions that may be created for a given basic block, and so on the number of entry points that may be created for any given function. If there are already too many specialized entry points for a given callee a generic entry point may be obtained instead. This does not matter to the caller and occurs rarely in practice.

Propagating types from callers to callees allows eliminating redundant type tests in the callee, but also makes it possible to pass arguments without boxing them, thereby reducing the overhead incurred by function calls. Note that our approach does not use any dynamic dispatch to propagate type information from callers to callees. It relies on information obtained from typed shapes to give us the identity of callees (both global functions and object methods) for free. With the current system, the identity of callees is known at compilation time in 97% of cases on average. When the identity of a callee is unknown, a generic entry point is used.

3.2 Call Continuation Specialization

Achieving full interprocedural type propagation demands passing the result type information from callees to callers. While it is fairly straightforward to establish the identity of the single callee a given call site will jump to in the majority of cases, the points at which a function returns to callers are likely to return to multiple call continuations within a program. These continuations may in turn receive input from multiple return points.

Type information about return values could be propagated with a dynamic dispatch of the return address indexed with the result type. However this would incur a run time cost. Instead, our second extension to BBV uses an approach with a zero run time cost amortized overhead. The call continuations are compiled lazily when the first return to a given con-

```
function sum(tree) {
    if (tree == null)
        return 0;
    else
        return sum(tree.left) +
               sum(tree.right) +
               t.val;
}

function makeTree(depth) {
    if (depth == 0)
        return null;
    else
        return { val: depth,
                 left: makeTree(depth-1),
                 right: makeTree(depth-1)
               };
}

var root = makeTree(8);
if (sum(root) != 502)
    throw Error('error');
```

Figure 4. Binary tree traversal example.

```
S1 = { val: int32, left: null,   right: null   }
S2 = { val: int32, left: object, right: object }
```

Figure 5. Typed shapes for tree nodes seen by the `sum` function

tinuation is executed. Each time a return statement is compiled, the type that is being returned is memorized (if known at compilation time).

When compiling a call continuation, the identity of the callee is checked. If the identity of the callee is known and its return type has been determined (so far), the code is specialized based on this type. If the callee, later during its execution, returns an incompatible type with the one seen so far, then all call sites of this function are notified that their call continuations are now invalid and must be recompiled. Continuations which are invalidated have their code replaced with a stub that will lazily trigger recompilation with an `unknown` result type if the continuation is ever executed again.

4. Extended Example

For illustration purposes, Figure 4 shows the `sum` function for traversing a balanced binary tree and computing the sum of numerical values stored in each node. Also shown are functions used to construct such a tree and test the behavior of the code. While this example may appear simple, there is much semantic complexity hiding behind the scene. A correct but naive implementation of this function contains many implicit dynamic tests.

Since our system uses typed shapes to specialize code based on object property types (see Section 2.3), the tree nodes produced by `makeTree` come in two flavors, shown in Figure 5. Shape `S1` corresponds to leaf nodes, and encodes that both `left` and `right` node pointers are set to null. Shape `S2` corresponds to internal nodes, and encodes that both node pointers are object references.

Figure 6 shows a diagram of the code generated for the `sum` function by Higgs with intraprocedural BBV. Some details, such as the details of global function lookups and overflow handling, were omitted from the figure for brevity. When entering the function, the type tag of `tree` must be tested. Either the value is `null` or it is an `object`. Once `tree` is known to be an `object`, its shape must be tested so that its `val`, `left` and `right` properties may be accessed. This is necessary because shapes encode the offsets in memory at which the properties are stored. The test for shape S2 is performed first because it is the shape of `tree` when that code was generated by BBV.

Whether `tree` is a leaf (shape S1) or internal node (shape S2), the code executed is abstractly the same. Two recursive calls to `sum` are made so that the sums can be computed for the `left` and `right` children of `tree`. Higgs knows that `tree.val` is an integer value because that information is encoded in `S1` and `S2`, but with intraprocedural BBV, does not know what the return types of the recursive calls to `sum` are. Hence, it must test that `t1` and `t2` are integers before adding them to `tree.val` with integer additions.

Figure 7 shows the code generated for the `sum` function with entry point versioning and continuation specialization enabled. Entry point versioning has created two versions of the `sum` function: one for when the tree is an `object` and one for when the tree is `null`. The version of `sum` specialized for null is trivial and returns immediately. The version specialized for objects does two shape checks, as before, but contains no type tag checks.

The elimination of type tag checks is possible because, on entry to `sum(tree:object)`, `tree` is known to be an object due to entry point specialization. Furthermore, continuation specialization determines that `sum` returns `int32` values, which eliminates the type tag checks on `t1`, `t2`, `t3`, and `t4`. The resulting code is faster and more compact.

5. Evaluation

An implementation of the Higgs JIT compiler extended with entry point and continuation specialization was tested on a total of 26 classic benchmarks from the SunSpider and V8 suites. One benchmark from the SunSpider suite and one from the V8 suite were not included in our tests because Higgs does not yet implement the required features. Benchmarks making use of regular expressions were discarded because Higgs and Truffle/JS [30, 31] do not implement JIT compilation of regular expressions.

To measure steady state execution time separately from compilation time in a manner compatible with both Higgs and Truffle/JS, the benchmarks were modified so that they

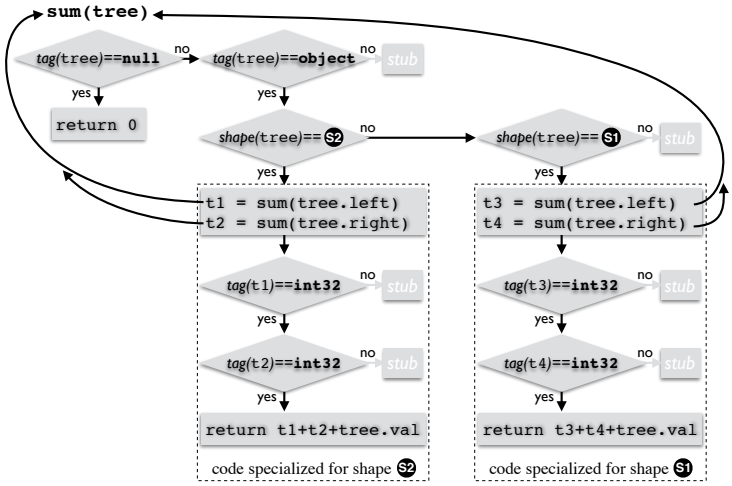

Figure 6. Generated code for the `sum` function with intraprocedural BBV

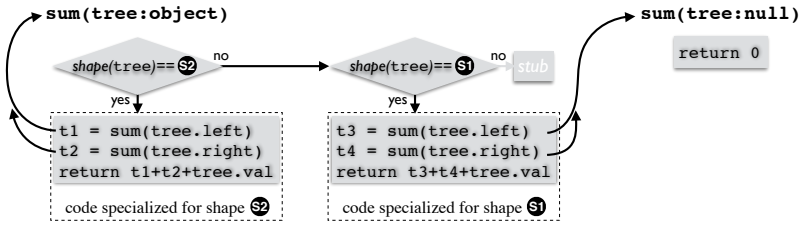

Figure 7. Generated code for the `sum` function with interprocedural BBV

could be run in a loop. A number of warmup iterations are first performed so as to trigger JIT compilation and optimization of code before timing runs take place.

The number of warmup and timing iterations were scaled so that short-running benchmarks would execute for at least 1000ms in total during both warmup and timing. Unless otherwise specified, all benchmarks were run for at least 10 warmup iterations and 10 timing iterations.

Truffle/JS v0.5 was used for performance comparisons. Tests were executed on a system equipped with an Intel Core i7-4771 CPU and 16GB of RAM running Ubuntu Linux 12.04. Dynamic CPU frequency scaling was disabled to ensure reliable timing measurements.

5.1 Dynamic Type Tests

It is useful to compare the performance of BBV in terms of type tag tests eliminated against a traditional static type analysis for JS. However, implementing such an analysis can be time consuming and error prone. It also raises the issue as to which specific kind of type analysis should be implemented, and whether the comparison is fair or not.

In order to avoid these issues, a simulated "perfect" analysis was used. Each benchmark was run and the result of all tag tests was recorded. The benchmarks were then rerun with all type tests that always evaluate to the same result removed. The simulated analysis is essentially an oracle with perfect knowledge of the future behavior of a program. It knows which type tests are redundant based on past executions on the same input. This gives an upper bound on the best performance achievable with a static type analysis in terms of tag tests eliminated with the constraint that the code is not transformed (for instance, no code duplication).

Figure 8 shows the relative proportion of type tests executed with different variants of BBV as well as with the simulated analysis. The first column shows the type tests remaining with plain intraprocedural BBV, as introduced in [12]. The second column shows the results obtained by adding typed shapes to intraprocedural BBV, as in [11]. The third column adds entry point specialization, and the fourth column finally adds call continuation specialization as introduced in this paper. Lastly, the fifth column shows the results of the perfect analysis.

Plain intraprocedural BBV eliminates 60% of type tests on average, and typed shapes brings us to 78%. The addition of entry point specialization improves the result further to 89% of type tests eliminated. Finally, completing our interprocedural implementation of BBV with the addition of call continuation specialization allows us to reach 94.3% of type tests eliminated, in several cases, nearly 100%.

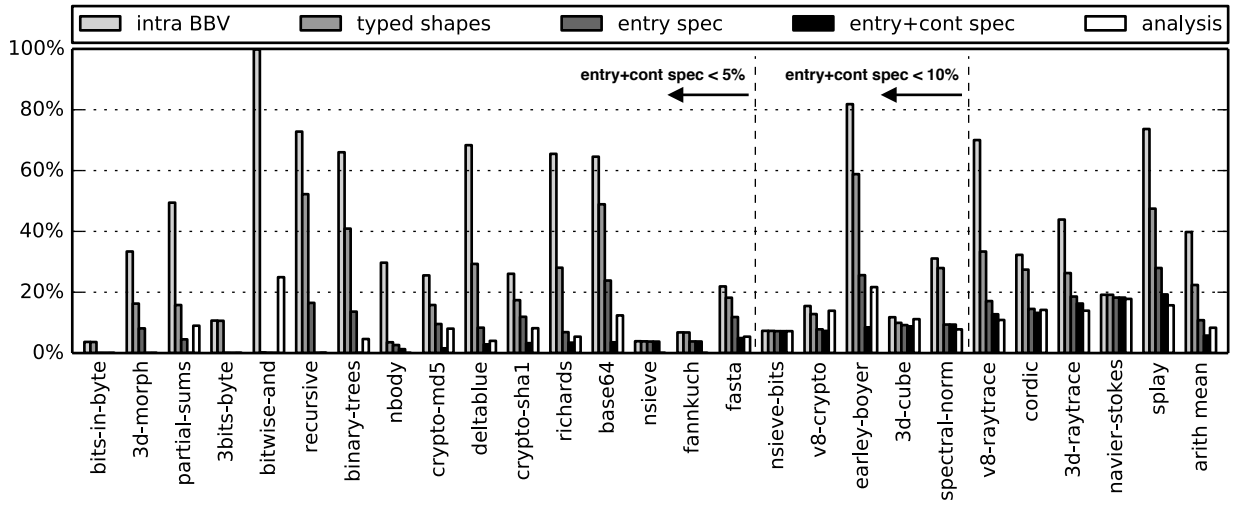

Figure 8. Proportion of type tests remaining relative to a baseline without BBV (lower is better)

It is surprising that interprocedural BBV performs as well or better than the perfect analysis on most benchmarks. The analysis eliminates just 91.7% of type tag tests on average, less than what interprocedural BBV achieves. BBV's ability to duplicate certain code paths allows it to track types more precisely and eliminate more type tests than is possible without transforming code.

5.2 Call Continuation Specialization

As outlined in Section 3.2, call continuation specialization uses a speculative strategy to propagate return type information without dynamic dispatch. That is, call continuations for a given callee function may be recompiled and deoptimized if values are returned which do not match previously encountered return types for the said function.

Empirically, across our benchmark sets, 10% of functions compiled cause the invalidation of call continuation code. Dynamically, the type tag of return values is successfully propagated and known to the caller 81% of the time. In over half of the benchmarks, the type tag of return values is known over 99% of the time.

5.3 Code Size

Adding entry point and continuation specialization to our existing BBV implementation cause an average increase in machine code size of 7.6% in the worst case and just 2.2% on average. Intuitively, one may have assumed that a bigger code size increase might have occurred, given that entry point versioning can generate multiple entry points per function. However, better optimized machine code tends to be more compact. We argue that the slight code size increase is observed is reasonable, particularly given the performance improvements obtained.

5.4 Compilation time

The addition of entry point and continuation specialization to our BBV implementation cause a compilation time increase of 4.7% in the worst case and 1.0% on average. The increase in compilation time is relatively small, and roughly corresponds to the increase in code size.

5.5 Execution Time

Figure 9 shows the relative execution time with entry point specialization as well as with both entry point and continuation specialization. When both entry point and continuation specialization are used, speedups of 14.5% are achieved on average, and up to 56% in the case of 3bits-bytes. Importantly, on nearly all benchmarks, the addition of continuation specialization does improve the average performance over entry point specialization only.

5.6 Comparison with Truffle/JS

Truffle/JS, like Higgs, is a research VM written in a garbage-collected language. For this reason it is interesting to compare their performance.

Since Truffle takes a relatively long time to compile and optimize code, each benchmark is run for 1000 "warm up" iterations before measuring execution time, so as to more accurately measure the final performance of the generated machine code once a steady state is reached. After the warm up period, each benchmark is run for 100 timing iterations.

Figure 10 shows the results of the performance comparison against Truffle/JS. The left column represents the speedup of Higgs over Truffle when taking only execution time into account, warmup iterations excluded. The right column is the speedup of Higgs over Truffle/JS when comparing total time, including initialization, compilation and warmup. A logarithmic scale was used due to the wide spread of data points.

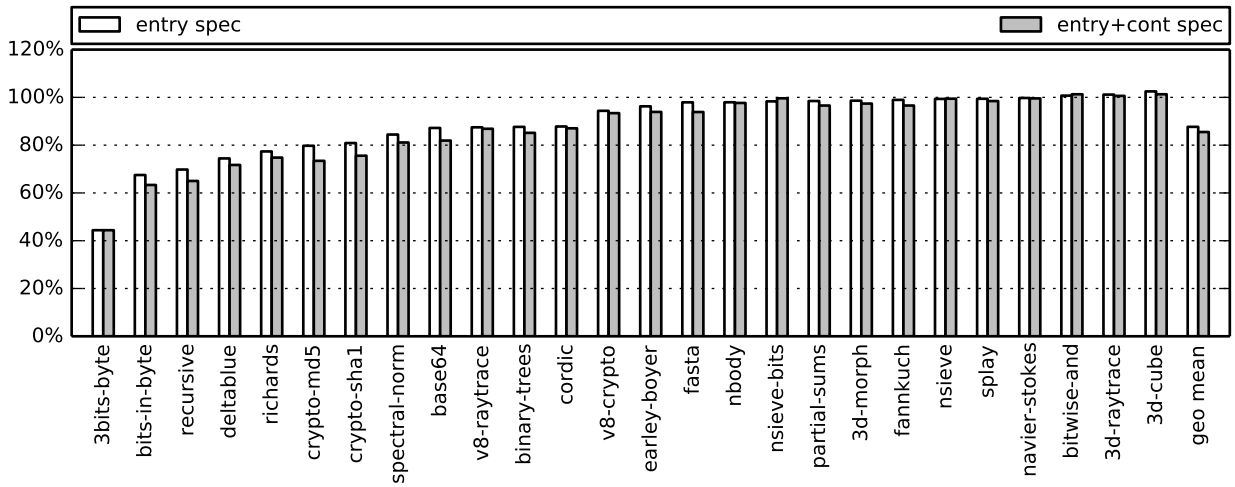

Figure 9. Execution time relative to baseline (lower is better)

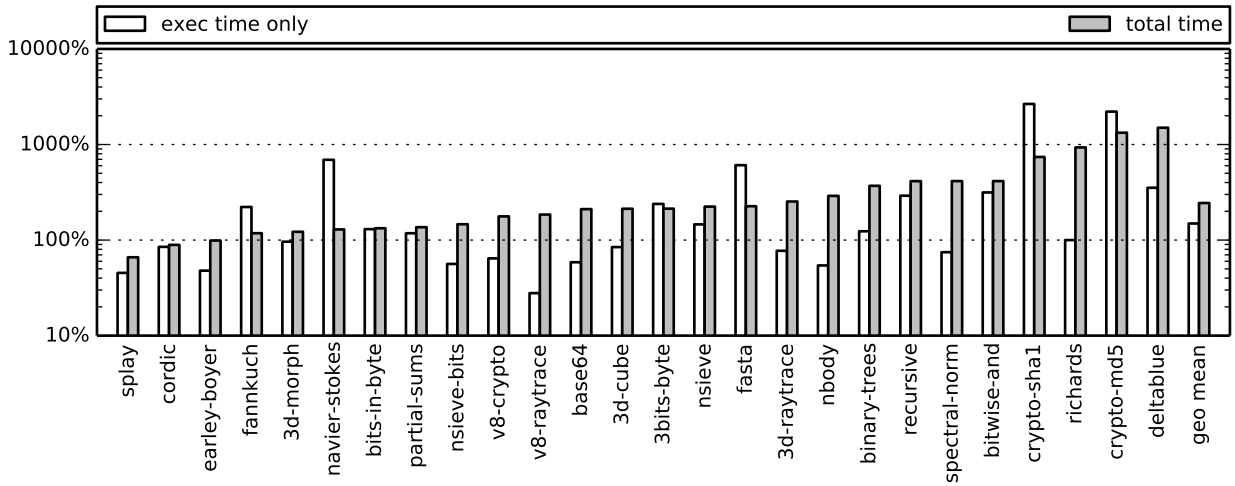

Figure 10. Speedup relative to Truffle (log scale, higher favors Higgs)

On average, when taking only execution time into account, Higgs outperforms Truffle on half of the benchmark set, and is 1.49× faster than Truffle/JS on average. When comparing wall clock times, Higgs outperforms Truffle on the majority of benchmarks, and is 2.44× faster on average.

It is not the case that Higgs outperforms Truffle/JS on every benchmark. JIT compiler design is a complex problem space, and as a result, there are benchmarks where each system outperforms the other by a wide margin. Truffle/JS is a method-based compiler implementing profiling, type-feedback and type analysis. It has a few technical advantages over Higgs, such as the use of method inlining, a more sophisticated register allocator, multithreaded background compilation, and a highly optimized garbage collector implementation (the Java VM's). These tools give Truffle an edge on specific benchmarks, such as `v8-raytrace`. Nevertheless, Higgs produces competitive machine code in many cases. It is interesting to see that using the simpler interprocedural BBV approach yields faster compilation (no warmup to speak of), and faster code on average. As a rough measure of compiler complexity, the source code of Truffle/JS is about 6 times larger than Higgs'.

6. Limitations and Future Work

While the results obtained with interprocedural BBV are in many cases superior to what is achievable with a traditional type analysis, there is still room for improvement. Our system does not keep track of the types of captured closure variables. It would be desirable to further extend BBV to do so. The system also treats arrays like untyped black boxes. Ex-

tending the system to keep track of the types of values stored inside of arrays, in particular numerical types, would likely lead to performance improvements.

7. Related Work

7.1 Type Analysis of Dynamic Languages

Basic block versioning bears resemblance to the *iterative type analysis* and *extended message splitting* techniques developed for Self by Craig Chambers and David Ungar [10]. This is a combined static analysis and transformation that compiles multiple versions of loops and duplicates control flow paths to eliminate type tests. The analysis works in an iterative fashion, transforming the control flow graph of a function while performing a type analysis. It integrates a mechanism to generate new versions of loops when needed, and a message splitting algorithm to try and minimize type information lost through control flow merges. One key disadvantage is that statically cloning code requires being conservative, generating potentially more code than necessary, as it is impossible to statically determine exactly which control flow paths will be taken at run time, and this must be overapproximated. Basic block versioning is simpler to implement and generates code lazily, requiring less compilation time and memory overhead, making it more suitable for integration into a baseline JIT compiler.

There have been multiple efforts to devise type analyses for dynamic languages. The Rapid Atomic Type Analysis (RATA) [24] is an intraprocedural flow-sensitive analysis based on abstract interpretation that aims to assign unique types to each variable inside of a function. Attempts have also been made to define formal semantics for a subset of dynamic languages such as JavaScript [4], Ruby [14] and Python [3], sidestepping some of the complexity of these languages and making them more amenable to traditional type inference techniques. There are also flow-based interprocedural type analyses for JavaScript based on sophisticated type lattices [20][21][22]. Such analyses are usable in the context of static code analysis, but take too long to execute to be usable in VMs and do not deal with the complexities of dynamic code loading.

The work done by Kedlaya, Roesch et al. [23] shows strategies for improving the precision of type analyses by combining them with type feedback and profiling. This strategy shows promise, but does not explicitly deal with object shapes and property types. Work has also been done on a flow-sensitive alias analysis for dynamic languages [17]. This analysis tries to strike a balance between precision and speed, it is likely too expensive for use in JIT compilers, however.

More recently, work done by Brian Hackett et al. at Mozilla resulted in an interprocedural hybrid type analysis for JavaScript suitable for use in production JIT compilers [19]. This analysis represents an important step forward for dynamic languages, but as with other type analyses, must conservatively assign one type to each value, making it vulnerable to imprecise type information polluting analysis results. Basic block versioning can help improve on the results of such an analysis by hoisting tests out of loops and generating multiple optimized code paths where appropriate.

7.2 Trace Compilation

Trace compilation, originally introduced by the Dynamo [5] native code optimization system, and later applied to JIT compilation in HotpathVM [16] aims to record long sequences of instructions executed inside hot loops. Such linear sequences of instructions often make optimization simpler. Type information can be accumulated along traces and used to specialize code and remove type tests [15], overflow checks [28] or unnecessary allocations [6]. Basic block versioning resembles tracing in that context updating works on essentially linear code fragments and code is optimized similarly to what may be done in a tracing JIT. Code is also compiled lazily, as needed, without compiling whole functions at once. Trace compilation [7] and meta-tracing are still an active area of research [8].

The simplicity of basic block versioning is one of its main advantages. It does not require external infrastructure such as an interpreter to execute code or record traces. Trace compiler implementations must deal with corner cases that do not appear with basic block versioning. With trace compilation, there is the potential for trace explosion if there is a large number of control flow paths going through a loop. It is also not obvious how many times a loop should be recorded or unrolled to maximize the elimination of type checks. This problem is solved with basic block versioning since versioning is driven by type information. Trace compilers must implement parameterizable policies and mechanisms to deal with recursion, nested loops and potentially very long traces that do not fit in instruction caches.

7.3 Just-In-Time Code Specialization

Customization is a technique developed to optimize Self programs [9] that compiles multiple copies of methods specialized on the receiver object type. Similarly, *type-directed cloning* [25] clones methods based on argument types, producing more specialized code using richer type information. The work of Chevalier-Boisvert et al. on *Just-in-time specialization* for MATLAB [13] and similar work done for the MaJIC MATLAB compiler [2] tries to capture argument types to dynamically compile optimized versions of whole functions. All of these techniques are forms of type-driven code duplication aimed at extracting type information. Basic block versioning operates at a lower level of granularity, allowing it to find optimization opportunities inside of method bodies by duplicating code paths.

There are notable similarities between the Psyco JIT specialization work and our own. The Psyco prototype for Python [26] is able to interleave execution and JIT compilation to gather run time information about values, so as to

specialize code on-the-fly based on types and values. It also incorporates a scheme where functions can have multiple entry points. We extend upon this work by combining a similar approach, that of basic block versioning, with typed shapes and a mechanism for propagating return types from callees to callers with low overhead.

The *tracelet-based* approach used by Facebook's *HHVM* for PHP [1] bears important similarities to our own. It is based on the JIT compilation of small code regions (tracelets) which are single-entry multiple-exit basic blocks. Each tracelet is type-specialized based on variable types observed at JIT compilation time. Guards are inserted at the entry of tracelets to verify at run time that the types observed are still valid for all future executions. High-level instructions in tracelets are specialized based on the guarded types. If these guards fail, new versions of tracelets are compiled based on different type assumptions and chained to the failing guards.

There are three important differences between the HHVM approach and basic block versioning. The first is that BBV does not insert dynamic guards but instead exposes and exploits the underlying type checks that are part of the definition of runtime primitives. HHVM cannot do this as it uses monolithic high-level instructions to represent PHP primitives, whereas the Higgs primitives are self-hosted and defined in an extended JavaScript dialect. The second difference is that BBV propagates known types to successors and doesn't usually need to re-check the types of local variables. Finally, HHVM uses an interpreter as fallback when too many tracelet versions are generated. Higgs falls back to generic basic block versions which do not make type assumptions but are still always JIT compiled for better performance.

8. Conclusion

Basic Block Versioning (BBV) is a JIT compilation strategy for generating type-specialized machine code on-the-fly without a separate type analysis pass. We have shown that it can be successfully extended with techniques to propagate information across method call boundaries, both from callers to callees and from callees to callers, and this without requiring dynamic dispatch.

Across 26 benchmarks, interprocedural BBV eliminates 94.3% of type tests on average, more than a simulated static type analysis with access to perfect information. The addition of interprocedural capabilities to BBV provides an additional speedup of 14.5% on average over an unextended BBV implementation.

Acknowledgments

Special thanks go to Laurie Hendren, Erick Lavoie, Vincent Foley, Paul Khuong, Molly Everett, Brett Fraley and all those who have contributed to the development of Higgs.

This work was supported, in part, by the Natural Sciences and Engineering Research Council of Canada (NSERC) and Mozilla Corporation.

References

[1] Keith Adams, Jason Evans, Bertrand Maher, Guilherme Ottoni, Andrew Paroski, Brett Simmers, Edwin Smith, and Owen Yamauchi. The hiphop virtual machine. In *Proceedings of the 2014 conference on Object Oriented Programming Systems Languages & Applications (OOPSLA)*, pages 777–790. ACM New York, 2014.

[2] George Almási and David Padua. MaJIC: compiling MATLAB for speed and responsiveness. In *Proceedings of the 2002 conference on Programming Language Design and Implementation (PLDI)*, pages 294–303. ACM New York, May 2002.

[3] Davide Ancona, Massimo Ancona, Antonio Cuni, and Nicholas D. Matsakis. RPython: a step towards reconciling dynamically and statically typed OO languages. In *Proceedings of the 2007 Dynamic Languages Symposium (DLS)*, pages 53–64. ACM New York, 2007.

[4] Christopher Anderson, Paola Giannini, and Sophia Drossopoulou. Towards type inference for JavaScript. In *Proceedings of ECOOP 2005*, pages 428–452. Springer Berlin Heidelberg, 2005.

[5] V. Bala, E. Duesterwald, and S. Banerjia. Dynamo: a transparent dynamic optimization system. In *Proceedings of the 2000 conference on Programming*, pages 1–12. ACM New York, 2000.

[6] Carl Friedrich Bolz, Antonio Cuni, Maciej FijaBkowski, Michael Leuschel, Samuele Pedroni, and Armin Rigo. Allocation removal by partial evaluation in a tracing JIT. In *Proceedings of the 20th ACM SIGPLAN workshop on Partial Evaluation and Program Manipulation (PEPM)*, pages 43–52. ACM New York, 2011.

[7] Carl Friedrich Bolz, Antonio Cuni, Maciej Fijalkowski, and Armin Rigo. Tracing the meta-level: Pypy's tracing jit compiler. In *Proceedings of the 4th workshop on the Implementation, Compilation, Optimization of Object-Oriented Languages and Programming Systems*, pages 18–25. ACM, 2009.

[8] Carl Friedrich Bolz, Tobias Pape, Jeremy Siek, and Sam Tobin-Hochstadt. Meta-tracing makes a fast Racket. *Workshop on Dynamic Languages and Applications*, 2014.

[9] Craig Chambers and David Ungar. Customization: optimizing compiler technology for Self, a dynamically-typed object-oriented programming language. In *Proceedings of the 1989 conference on Programming Language Design and Implementation (PLDI)*, pages 146–160. ACM New York, June 1989.

[10] Craig Chambers and David Ungar. Iterative type analysis and extended message splitting: Optimizing dynamically-typed object-oriented programs. In *Proceedings of the 1990 conference on Programming Language Design and Implementation (PLDI)*, pages 150–164. ACM New York, 1990.

[11] Maxime Chevalier-Boisvert and Marc Feeley. Extending basic block versioning with typed object shapes. Technical report,

2015. http://arxiv.org/abs/1507.02437.

[12] Maxime Chevalier-Boisvert and Marc Feeley. Simple and effective type check removal through lazy basic block versioning. In *29th European Conference on Object-Oriented Programming (ECOOP 2015)*, volume 37 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 101–123. Schloss Dagstuhl, 2015. http://arxiv.org/abs/1411.0352.

[13] Maxime Chevalier-Boisvert, Laurie Hendren, and Clark Verbrugge. Optimizing MATLAB through just-in-time specialization. In *Proceedings of the 2010 international conference on Compiler Construction (CC)*, pages 46–65. Springer Berlin Heidelberg, 2010.

[14] Michael Furr, Jong-hoon (David) An, Jeffrey S. Foster, and Michael Hicks. Static type inference for ruby. In *Proceedings of the 2009 ACM Symposium on Applied Computing (SAC)*, pages 1859–1866. ACM New York, 2009.

[15] Andreas Gal, Brendan Eich, Mike Shaver, David Anderson, David Mandelin, Mohammad R. Haghighat, Blake Kaplan, Graydon Hoare, Boris Zbarsky, Jason Orendorff, Jesse Ruderman, Edwin W. Smith, Rick Reitmaier, Michael Bebenita, Mason Chang, and Michael Franz. Trace-based just-in-time type specialization for dynamic languages. *SIGPLAN Not.*, 44(6):465–478, June 2009.

[16] Andreas Gal, Christian W. Probst, and Michael Franz. HotpathVM: an effective JIT compiler for resource-constrained devices. In *Proceedings of the 2nd international conference on Virtual Execution Environments (VEE)*, pages 144–153. ACM New York, 2006.

[17] Michael Gorbovitski, Yanhong A. Liu, Scott D. Stoller, Tom Rothamel, and Tuncay K. Tekle. Alias analysis for optimization of dynamic languages. In *Proceedings of the 6th Symposium on Dynamic Languages*, DLS '10, pages 27–42, New York, NY, USA, 2010. ACM.

[18] David Gudeman. Representing type information in dynamically typed languages, 1993.

[19] Brian Hackett and Shu-yu Guo. Fast and precise hybrid type inference for JavaScript. In *Proceedings of the 33rd ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI)*, pages 239–250. ACM New York, June 2012.

[20] Simon Holm Jensen, Anders Møller, and Peter Thiemann. Type analysis for JavaScript. In *Proceedings of the 16th International Symposium on Static Analysis (SAS)*, pages 238–255. Springer Berlin Heidelberg, 2009.

[21] Simon Holm Jensen, Anders Møller, and Peter Thiemann. Interprocedural analysis with lazy propagation. In *Proceedings 17th International Static Analysis Symposium (SAS)*. Springer Berlin Heidelberg, September 2010.

[22] Vineeth Kashyap, John Sarracino, John Wagner, Ben Wiedermann, and Ben Hardekopf. Type refinement for static analysis of JavaScript. In *Proceedings of the 2013 Dynamic Languages Symposium (DLS*. ACM New York, 2013.

[23] Madhukar N. Kedlaya, Jared Roesch, Behnam Robatmili, Mehrdad Reshadi, and Ben Hardekopf. Improved type specialization for dynamic scripting languages. *SIGPLAN Not.*, 49(2):37–48, October 2013.

[24] Francesco Logozzo and Herman Venter. RATA: rapid atomic type analysis by abstract interpretation; application to JavaScript optimization. In *Proceedings of the 2010 international conference on Compiler Construction (CC)*, pages 66–83. Springer Berlin Heidelberg, 2010.

[25] John Plevyak and Andrew A. Chien. Type directed cloning for object-oriented programs. In *Proceedings of the Workshop for Languages and Compilers for Parallel Computing (LCPC)*, pages 566–580, 1995.

[26] Armin Rigo. Representation-based just-in-time specialization and the psyco prototype for python. In *Proceedings of the 2004 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation*, PEPM '04, pages 15–26, New York, NY, USA, 2004. ACM.

[27] Baptiste Saleil and Marc Feeley. Code versioning and extremely lazy compilation of scheme. In *Scheme and Functional Programming Workshop*, 2014.

[28] Rodrigo Sol, Christophe Guillon, FernandoMagno Quinto Pereira, and Mariza A.S. Bigonha. Dynamic elimination of overflow tests in a trace compiler. In Jens Knoop, editor, *Proceedings of the 2011 international conference on Compiler Construction (CC)*, pages 2–21. Springer Berlin Heidelberg, 2011.

[29] Andreas Wöß, Christian Wirth, Daniele Bonetta, Chris Seaton, Christian Humer, and Hanspeter Mössenböck. An object storage model for the truffle language implementation framework. In *Proceedings of the 2014 International Conference on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools*, pages 133–144, New York, NY, USA, 2014. ACM.

[30] Thomas Würthinger, Christian Wimmer, Andreas Wöß, Lukas Stadler, Gilles Duboscq, Christian Humer, Gregor Richards, Doug Simon, and Mario Wolczko. One VM to rule them all. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software*, Onward! 2013, pages 187–204. ACM New York, 2013.

[31] Thomas Würthinger, Andreas Wöß, Lukas Stadler, Gilles Duboscq, Doug Simon, and Christian Wimmer. Self-optimizing AST interpreters. In *Proceedings of the 2012 Dynamic Language Symposium (DLS)*, pages 73–82. ACM New York, 2012.